

COLOR VISION FOR ROBOTIC ORANGE HARVESTING

By

DAVID CHARLES SLAUGHTER

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN
PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

1987

ACKNOWLEDGMENTS

The author wishes to express gratitude for the guidance and assistance his major advisor, Dr. R.C. Harrell, provided throughout the author's graduate program. The support of K. Norris of the Sensors and Control Systems Institute, USDA-ARS, and Dr. J.C. Webb of the Insect Attractants and Basic Biology Research Laboratory, USDA-ARS, were an invaluable asset to this research. The assistance of Dr. J.D. Whitney in helping to arrange for the filming of orange groves was greatly appreciated. The author would also like to show his appreciation to his wife, Susan, for her patient support and assistance.

TABLE OF CONTENTS

	PAGE
ACKNOWLEDGMENTS	ii
ABSTRACT	v
INTRODUCTION	1
The Citrus Harvest Problem	1
Objectives	3
LITERATURE REVIEW	6
Challenges	6
The 'Valencia' Cultivar: A Challenge to Mechanical Harvest	6
The Challenge to Robotic Orange Harvest	7
Robotic Tree Fruit Harvesting	8
Vision Sensing	11
Machine Vision	11
Real-Time Vision Requirements	13
Spectral Reflectance Information	14
Image Enhancement Using Interference Filters ..	17
The Color Video Camera	18
Color Vision	20
Color Theory	20
Color Video Information	24
Encoding color information	24
Decoding color information	25
Color Machine Vision	26
PROCEDURE	32
Overview of Research	32
Quantifying Color Information in Natural Orange Grove Scenes	32
Color Imaging	34
Hue and saturation thresholding	34
Statistical pattern classification	35
Quantifying color segmented image quality	36
Real-Time Vision for Robotic Guidance	38
Aperture Control	41
Image quality and illumination	41
A new measure of image intensity	42

	PAGE
Equipment Overview	42
Robotic Fruit Harvester with Color Vision System	42
Color Vision System with Object Oriented Aperture Control	45
System Hardware	45
Spectrophotometric hardware	45
Color vision hardware	46
Aperture control hardware	48
Implementation	48
Determining the Colors of a Typical Orange Grove Scene	48
Processing Color Images	50
Taping natural orange grove scenes	50
Developing color lookup tables	51
Statistical pattern classification	54
Measuring color segmented image quality	56
Evaluating color segmentation	59
Evaluating Real-Time Orange Location	62
Aperture Control Experiments	63
Lens aperture--image intensity relationships	63
Characterizing autoiris capabilities	64
RESULTS	68
Color Characteristics of Typical Orange Grove Objects	68
Color Segmentation of Natural Orange Grove Scenes	77
Real-Time Considerations	94
Aperture Control	95
CONCLUSIONS	107
REFERENCES	112
BIOGRAPHICAL SKETCH	117

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

COLOR VISION FOR ROBOTIC ORANGE HARVESTING

By

DAVID CHARLES SLAUGHTER

August 1987

Chairman: R.C. Harrell
Major Department: Agricultural Engineering

Color vision was investigated as a means of providing real-time guidance information for the control of a robotic orange harvester. Detecting and locating a fruit by its color rather than its shape or gray level greatly reduced the complexity of the problem.

This study focused on four major issues. First, what were the color characteristics of typical objects in natural orange grove scenes? Second, how could color be used to detect and locate oranges? Third, once a color vision algorithm was developed, what was its suitability for real-time robot guidance? Fourth, could adequate illumination control be provided using traditional autoiris hardware?

The diffuse spectral reflectance (visible spectrum) of typical orange grove objects was studied. The natural

contrast in color between oranges and orange grove background objects was most obvious when the spectral reflectance information was translated into the intensity, hue and saturation color coordinate system.

A multivariate statistical classification technique was used to systematically classify pixels (picture elements) from a natural orange grove scene as either orange or background. This technique performed equally well if the color information was specified in terms of its red, green and blue components or its hue and saturation levels.

A real-time search algorithm was implemented in conjunction with a color lookup table for pixel classification. In the worst case, a fruit could be detected, its centroid and diameter estimated in 10.8 ms. The estimated centroid differed, on the average, from the true centroid by $\pm 10\%$ of the diameter of the fruit.

Quality of color segmented images was optimum when the average intensity of orange pixels was in the middle of its dynamic range. An object oriented aperture control system, that controlled the average intensity of the orange pixels, could maximize image quality. The dynamic response of a typical autoiris lens was too slow to respond to variations in illumination encountered when robot arm rapidly extends into the canopy of an orange tree.

Although oranges were studied in this work, the ideas presented apply equally well to most fruits that differ in color from the foliage.

INTRODUCTION

The Citrus Harvest Problem

The harvest of citrus is a labor intensive task requiring large numbers of workers for only a few months out of the year. The citrus farmer must recruit large numbers of employees who are willing to work for a short time harvesting citrus and then who must look for another source of employment. Martin (1983) found that three major factors contribute to the problem of fruit harvest. First, the ratio between high and low work force requirements in fruit harvesting is as much as 20:1. Second, the cost of hand harvesting citrus is 20 percent of the price the farmer gets for oranges and lemons. Third, wages of farm laborers are 5 times higher in the U.S. than they are in Greece and 10 times higher than they are in Mexico. Not only does the citrus farmer have a potential labor shortage (already large numbers of illegal aliens are thought to be working in the orange groves) but, even when the quantity of laborers is sufficient, the high cost of labor makes it difficult to be competitive on the world market. One possible solution to this problem is to mechanize fruit harvest in order to reduce the need for high volumes of seasonal laborers.

Of all operations involved in orange production, harvesting the fruit is the only operation still not mechanized. Coppock (1977) observed that cultivating, fertilizing, topping, hedging, spraying, sizing fruit, packing fruit and extracting juice have all been mechanized to some extent, but at the same time over 30 years of formal research have failed to produce a mechanical harvesting system that has received large scale industry acceptance. Although many mechanical harvesting systems for processed fruit have been developed, their feasibility under existing conditions has not been demonstrated. Some of the reasons for the lack of acceptance of harvest mechanization have been a high initial capital outlay, inefficient fruit recovery, and fear of permanent damage to trees.

Increasingly, our society looks to technological progress to provide high quality food at a minimal cost. Pejisa and Orrock (1983) suggested that citrus harvest was a likely candidate for an intelligent robotic harvesting system based upon total US farm gate value, crop value/acre, and manpower requirements. Before much progress can be made in this area, new sensor technology must be developed. Fruit location information is required before the task of robotic harvesting can be implemented. A high level of accuracy in determining fruit location is required because each step in the task of fruit harvest builds upon the guidance information. The development of new sensory capabilities will allow robotics to advance beyond the

preprogrammed environment of an industrial manipulator to the dynamic environment of agriculture. The scope of this research addresses the development of a sensing system for the location of oranges in their natural environment, the orange grove.

Objectives

The overall objective of this research was to investigate the feasibility of using a color vision system to provide guidance information for the control of a robotic manipulator during orange harvest. The specific objectives of this research were

- I To quantify the color characteristics of a natural orange grove scene.
- II To develop a technique using color information for detecting and locating oranges in natural orange grove scenes.
- III To evaluate the real-time suitability of implementing color image processing algorithms in a robotic fruit harvesting application.
- IV To investigate the feasibility of aperture control for locating oranges under varying illumination.

The first objective was accomplished using standard diffuse spectral reflectance techniques. The spectral information was translated into three commonly used color systems. This information provided an estimate of the potential for the use of color as a means of detecting and locating oranges in a natural orange grove scene. In addition the three color systems were examined for their potential in a color vision system.

The second objective was accomplished through the application of image processing and pattern classification techniques to the color information present in natural orange grove scenes. The scope of this objective was restricted to locating only the fruit that was visible from the exterior of the tree, and no attempt was made to differentiate single fruit from clustered fruit. Previous research indicated that these restrictions in the scope of this objective were appropriate. Schertz and Brown (1968) estimated that from 70% to 100% of the fruit in an orange tree was observable from the exterior of the tree. Brown et al. (1971) after studying 'Valencia' oranges in three counties in California determined that 70% of the fruit occur singly and 20% occur in clusters of two. The evaluation of the results for this objective were designed to be consistent with the overall objective of providing guidance information for the control of a robotic fruit harvester.

The third objective was accomplished through the analysis of the natural frequency of oscillation of oranges. Once the term real-time had been defined for this system, the feasibility of implementing the detection and location of oranges in a color image in real-time was investigated.

The fourth objective was accomplished through the analysis of the relationship between the quality of a color image and the changes in the color information in that image with varying illumination. The dynamic response of a typical autoiris lens was studied to assess its real-time suitability. The effects upon scene illumination through the use of artificial illumination were beyond the scope of this objective.

LITERATURE REVIEW

This chapter begins with a brief introduction to some of the challenges facing the development of a robotic orange harvesting system. Recent developments in robotic fruit harvesting are described. The need for better sensing systems for locating fruit in an agricultural environment is observed and several of the limitations in implementing a real-time vision control system with current machine vision technology are demonstrated. The chapter concludes with an overview of color theory and a few examples of the application of color vision to other fields.

Challenges

The 'Valencia' Cultivar: A Challenge to Mechanical Harvest

'Valencia' is a commonly grown cultivar of orange in the state of Florida having the unique trait of requiring fifteen months time from bloom to harvest. While grown for its reputation as a good juice orange, the 'Valencia' has caused problems for mechanical harvesters because of the presence of small, green immature fruit on the tree during the harvest of the mature fruit. This phenomenon challenges mechanical harvesting systems to successfully harvest the mature fruit while leaving the fruit for next year's crop on

the tree. Whitney (1977) reported that when removal rates of 85% to 90% of mature 'Valencia' fruit were achieved using mass shake harvesting techniques, the next year's crop was reduced by 15% to 20% due to removal of immature fruit. A robotic fruit harvesting system has the potential for meeting the challenge of selectively harvesting fruit like the human harvester, minimizing any reduction in next year's crop.

The Challenge to Robotic Orange Harvest

Guiding a robotic manipulator, from the initial detection and location of an orange in the canopy of a tree to the successful harvest and safe storage of the fruit, is not a simple task especially if the whole process is to be conducted at rates faster than that of a human picker. Harmon (1982), after considering the currently available robotics technology, was not optimistic about the application of robotics to fruit harvest in such a manner that could afford the human picker worthwhile competition in the orange grove. But at the same time the job of harvesting citrus is shunned by almost anyone who can find some other kind of labor paying the same wage. Coppock and Jutras (1960) reported that on the average, a hand fruit harvester picks about 40 fruit per minute when actively picking and spends only 75% of the work day actively picking fruit; the other 25% of the time is spent positioning ladders and transporting fruit.

Harrell (in press) determined that a robotic harvester could compete on an economic basis with traditional harvest methods if the robotic harvesting system could pick at least 93% of the oranges on the tree. A higher level of harvest inefficiency could be competitive if there was a corresponding increase in the cost of hand labor. Although this high level of performance is not easily attainable with current technology, it is not impossible. Harrell concluded that, as more research is conducted coupled with a decrease in the cost of robotic technology and a probable increase in labor costs, robotic harvesting of citrus is likely to become viable.

Robotic Tree Fruit Harvesting

In some of the earliest research in robotic fruit harvesting Parrish and Goksel (1977) demonstrated the technical feasibility of using machine vision to guide a spherical coordinate (RRP) robot in apple harvesting. An RRP robot has three degrees of freedom implemented with two rotational (R) joints and one prismatic (P) or sliding joint. In this research a standard black-and-white television camera was used to detect and locate apples. A color filter was used (in front of the camera lens) to enhance the contrast of fruit against background and to decrease the effects of intensity variations caused by illumination gradients or shadows. Although the system investigated was quite rudimentary in nature and never

picked a single apple, the results indicated the feasibility of a machine vision guided manipulator for fruit harvest and have provided a basis on which other researchers have built.

Grand d'Esnon (1984 and 1985) also investigated robotic apple harvesting. A CCD (charge couple device) line scan camera with optical interference filters was used to locate the horizontal coordinate of the apple to be picked; the vertical coordinate of the apple was known by the position and orientation of the camera. Dead reckoning guidance was used once the two dimensional location of the fruit was established and a photosensitive emitter/detector pair was used to determine when the end-effector was close enough to the fruit to pick it. A cylindrical coordinate (PRP) robot was used with the optical axis of the camera mounted parallel to the direction of travel of the picking arm. Using only one optical filter the vision system could detect fruit against foliage under cloudy conditions or at night with artificial illumination, but two or three optical filters were thought to be necessary to locate fruit in the sunshine. The system, with no a priori knowledge about the location of fruit on the tree, could harvest apples at a rate of approximately 15 fruit per minute.

A robotic system that would harvest citrus at night was proposed by Tuttle (1983). Tuttle proposed to use a photosensitive array with appropriate optical filters to guide the robot based on the ratio of light reflected from the scene in the 600 to 700 nm spectral region to that

reflected in the 750 to 850 nm region. This imaging scheme was proposed to compensate for the fact that the energy reflected from a surface is inversely proportional to the distance from the surface raised to the fourth power. If a single optical filter was used, as in the research done by Parrish and Goksel (1977), a leaf 1 m from the image sensor could theoretically appear brighter than an orange 3 m from the image sensor, confusing the system. Night harvest was required because oranges in the shade do not necessarily reflect more light than leaves in the sun (Schertz and Brown, 1968).

Harrell et al. (1985) investigated the use of real-time vision servoing of a robotic manipulator to harvest oranges. This system used a small black-and-white CCD camera mounted in the end-effector so that the optical axis of the camera was co-axial with the prismatic joint of the spherical coordinate (RRP) robot used. By mounting the camera in the end-effector rather than at the back of the robot the calculations involved in determining the location of the fruit relative to the end-effector were greatly simplified. Under simulated night harvest conditions, a high contrast image of plastic fruit and plastic foliage against a black background was obtained by using a red color filter in front of the camera lens. A gray level threshold was applied to segment the image into fruit and background regions. Once segmented a spiral search was performed starting in the center of the image and any object in the image satisfying a

minimum size (diameter) criterion was classified as an orange. The vision system provided two dimensional information to the vision-servo control routine at standard television frame rates (30 Hz). The distance to the fruit was estimated from its horizontal diameter in the image because plastic fruit, all of the same size and shape, were used to evaluate the system. This system was capable of harvesting plastic fruit from a simulated plastic orange tree at a rate of 15 fruit per minute.

The research conducted to date has attempted to show the technical feasibility of a robotic fruit harvesting system. Although progress toward this goal has been made, the challenge still exists for the development of a harvesting system that can out perform traditional harvest methods. One of the main areas of research that needs further attention is in sensor development. More research on the detection and location of a fruit in the natural environment of an orange grove must be made before a robotic harvesting system can expect to meet the challenge of orange harvest.

Vision Sensing

Machine Vision

New methods of detecting and locating objects in three dimensions must be developed before progress can be made toward robotic harvesting. The science of machine vision can be thought of as the process of recovering

three-dimensional information from a two-dimensional image. Compared with human vision, machine vision is unrefined. In comparison with the human vision system, which does the task of 100 low level operations simultaneously and has a frame rate of about 10 frames per second, Moravec (1984) concluded that a computer vision system, processing one million instructions per second, is about 10,000 times too slow to perfectly mimic the human vision system. Advanced cameras exist that have 1,000,000 picture elements (pixels) whereas the human eye has 250,000,000 sensing elements (Hackwood and Beni, 1984). It is the ability to readily interpret sight that allows the laborer to locate fruit in the foliage of a citrus tree.

Many computer vision techniques have been developed to detect and locate objects in three dimensions. Although these complex techniques, especially those involving artificial intelligence concepts, often produce impressive results, they are impractical for use in mobile robotic systems requiring real-time sensory feedback. For example, Katsushi et al. (1984) implemented a vision task of finding a target object and determining a grasping point using photometric stereo and a proximity sensor. The vision portion of the task required 40 to 50 seconds to acquire and process using a Lisp machine. Nakagawa and Ninomiya (1984) developed a vision system capable of detecting solder joints in 0.1 seconds but required structured lighting. Whittaker et al. (1984) used the circular Hough transform technique to

locate the center of tomatoes in natural scenes and Wolfe and Swaminathan (1986) used the circular Hough transform to identify bell peppers. In both of these applications of the circular Hough transform the technique was chosen for its robustness in that it is valid for partially occluded circular objects and works well even if the object is not perfectly circular. Unfortunately, high quality results depended upon the image being preprocessed by a Sobel type operation and the computational complexity of the entire process makes real-time application on a typical microprocessor based system unfeasible. Requirements such as carefully constrained environments (such as those only possible in the laboratory), massive computational power or the large amounts of processing time required to implement these techniques on a typical microprocessor make these techniques unfeasible for use in the control strategy of a robotic harvester.

Real-Time Vision Requirements

Real-time digital control of a robotic manipulator requires a feedback signal at a high sampling rate (50 Hz or higher) in order to achieve the dynamic performance necessary to harvest potentially swaying fruit. Due to this constraint it is not practical to consider many traditional pattern recognition techniques to find oranges in a standard video image. One technique commonly used to locate objects in an image is called image segmentation (Rosenfeld and Kak,

1982). The purpose of image segmentation is to divide the image into meaningful regions. The simplest form of segmentation is a binary image, an image with only two distinct regions (in this case fruit and background). One of the fastest methods of acquiring a binary image is to use gray level thresholding. Gray level thresholding requires that objects and background in the image have unique levels of brightness. The threshold is the brightness level that allows objects to be discriminated from the background. Segmentation, using gray level thresholding, can be performed extremely fast since the operation is easily handled in hardware at standard video rates. Once a binary image has been constructed, a quick and simple Boolean operator is sufficient to determine if a pixel is object or background. The main difficulty with gray level thresholding lies in the ability to choose a threshold value that adequately distinguishes object from background. In the case of oranges, the natural illumination in an orange grove is such that it is not known a priori whether the orange is brighter than the background or visa versa and by how much.

Spectral Reflectance Information

To overcome some of these problems researchers have searched for naturally present features of an orange in an orange tree that would help simplify the complex task of locating a fruit in among the foliage. Schertz and Brown

(1968) suggested that location of fruit might be accomplished by photometric information, specifically by using the light reflectance differences between leaves and fruit in the visible or infrared portion of the electromagnetic spectrum. Gaffney (1969) determined that 'Valencia' oranges could be sorted by color using a single wavelength band of reflected light at 660 nm. This technique was capable of distinguishing between normal orange, light orange, and regreened fruit. Coppock (1983) considered color as a possible criterion for locating citrus fruit in the tree but did not pursue the concept for lack of an effective system for sensing the color at each location in the tree effectively.

The spectral reflectance curves, in the visible region (400 nm to 700 nm), show a large difference (approximately 10 to 1 at 675 nm) between the amount of light reflected from the peel of an orange and the leaf of an orange tree (Figure 1). This difference is due primarily to the presence of chlorophyll in the leaf which has a strong absorption band centered at 675 nm (Hollaender, 1956). As a result of the difference in spectral reflectance characteristics, light from 600 nm to 700 nm (as when viewed through a red interference filter) allows a vision system to distinguish between fruit and leaves using only brightness information. The spectral reflectance information (Figure 1) was plotted as the logarithm of the reflectance, R , to accentuate differences between the spectra.




Figure 1. Spectral reflectance curves for orange peel and orange tree leaf.

Unfortunately, fruit and background are not so easily differentiated in the grove. Background sky, clouds, and soil often have high reflectances in the 600 nm to 700 nm region of the spectrum causing difficulties for a segmentation process based entirely upon brightness. In addition, intensity differences as great as 40 to 1 under natural illumination in the canopy of an orange tree have been measured (Schertz and Brown, 1968). Thus, an orange in the sun would be ten times brighter than a leaf in the sun, while a leaf in the sun could appear four times brighter than an orange in the shade.

Image Enhancement Using Interference Filters

The key to successful image segmentation is distinct and non-overlapping levels of brightness between object and background in the original image. In the laboratory, a narrow band pass filter centered at 680 nm can be used to distinguish fruit from leaves, but in the grove this method could misclassify sky, clouds, and soil as fruit. Night harvest might be possible using this method with structured lighting, although soil reflectance would still be a problem (Tuttle, 1983).

One method of attacking this problem would be to subtract two images taken of the same scene with different interference filters. The reflectance of citrus fruit in the region from 410 nm to 480 nm is very low (5%), whereas background sky, clouds, and sandy soil have uniformly high

illumination in the visible spectrum. Slaughter et al. (1986) demonstrated that the image resulting from subtraction of an image filtered at 450 nm from an image of the same scene filtered at 680 nm can be segmented to classify fruit, trees, and sky correctly (Figure 2). A major disadvantage to this method is the requirement that two images of the same scene must be taken using two different narrow-band-pass filters (requiring two separate cameras or mechanically switching filters). Any spatial offset between images due to movement of the fruit or tree, or offset between cameras, complicates the process considerably especially if the fruit are partially occluded from view by leaves. In addition, any non-fruit object having high reflectance at 680 nm and low reflectance at 450 nm would be misclassified as fruit.

The Color Video Camera

The use of a color video camera greatly simplifies this problem by acquiring three optically filtered images of the same scene simultaneously. Technological advances have produced solid-state color video cameras which, in addition to their small size and low cost, are well suited to the task of searching for fruit in orchards because their sensors are not permanently damaged by intense illumination. Information from the three filtered images can be used to segment a natural scene into object and background regions based upon true color information.



(a)



(b)



(c)

Figure 2.

Black and white images of citrus tree with two fruit.

(a) Optically filtered with a narrow band pass filter centered at 680 nm.

(b) Optically filtered with a narrow band pass filter centered at 450 nm.

(c) Results from subtracting (b) from (a).

Color Vision

Color Theory

Color is often thought to be a property associated with particular objects; however, a more appropriate view is that color is a property of light. An object's color comes from the interaction of light waves with electrons in the object matter (Nassau, 1980). Thus an object has color only in the sense that it has the ability to modify the color of the light incident upon it. From an engineering standpoint visible light consists of a small region of electromagnetic radiation from 380 nm to 780 nm in the wavelength domain. Light, as defined by the Committee on Colorimetry of the Optical Society of America, is "the aspect of radiant energy of which a human observer is aware through the visual sensations which arise from the stimulation of the retina of the eye" (1944, p. 245). The Committee on Colorimetry defines color as "the characteristics of light other than spatial and temporal inhomogeneities" (p. 246).

By the seventeenth century a considerable amount was known about the properties of light, but little was known about color. The first steps toward understanding color were made by Issac Newton (1730). Newton found that the spectrum of colors, created by passing sunlight through a glass prism, could be combined back into "white" sunlight again by passing the color spectrum through a second inverted glass prism. Maxwell's triangle (an equilateral triangle named after J.C. Maxwell who also researched color

theory) is often used to represent the stimuli of additive combinations of three colored lights. The vertices of the triangle represented the three colored lights to be studied and were called primaries. Although any set of different colors can be used as primary colors, there are two commonly used sets of primary colors, termed additive and subtractive primaries. The additive primaries are red (R), green (G), and blue (B), whereas the subtractive primaries are yellow, cyan and magenta (Overheim and Wagner, 1982). The subtractive primaries (or pigment primaries) are used in the printing process while the additive primaries are used for combining sources of illumination and are the primaries used for video imaging. Most naturally occurring colors can be represented by an additive combination of three primary colors with the most notable exception being monochromatic light (e.g. a sodium flame).

Grassman (1853) developed several laws of color that became the basis for later work in colorimetry. Grassman's first law states that the perception of color is tridimensional or that the human eye is sensitive to three properties, luminance, dominant wavelength and purity (also known as brightness (Y), hue (θ) and saturation (P) respectively). The properties of brightness, hue and saturation are psychological properties not psychophysical properties. For example, red, orange, yellow, green, blue, and violet are some commonly used hues. Saturation relates to the strength of the hue and the terms deep, vivid, pale,

and pastel are examples of terms used to describe the saturation of a color. Brightness pertains to the intensity of the stimulation.

The International Commission on Illumination (CIE) developed a quantitative system for describing color (Judd, 1933). The CIE system is based upon three tristimulus values (or imaginary primaries) X, Y, and Z and is a precise refinement of Maxwell's color triangle. The design of the XYZ system was based upon imaginary primaries rather than a system based on real primaries (e.g. RGB) so that any color could be matched without requiring a mixture of negative intensities of the primaries. Further the Y primary was chosen to represent all of the luminosity (photometric brightness) of the color being matched. There is a unique relationship between each color and its triplet of XYZ values which enables the CIE system to be a standardized method for describing color as perceived by the human vision system.

The CIE tristimulus values can be calculated from spectral reflectance data using the following equations (Driscoll and Vaughan, 1978):

$$X = k \sum_{\lambda=380}^{780} \bar{x}(\lambda) \bar{g}(\lambda) \Delta\lambda , \quad (1)$$

$$Y = k \sum_{\lambda=380}^{780} \bar{y}(\lambda) \bar{g}(\lambda) \Delta\lambda , \quad (2)$$

$$Z = k \sum_{\lambda=380}^{780} \sigma(\lambda) \bar{z}(\lambda) \Delta\lambda, \quad (3)$$

where λ is the wavelength in nm. The color-stimulus function, $\sigma(\lambda)$, is determined by

$$\sigma(\lambda) = \rho(\lambda) S(\lambda), \quad (4)$$

where $\rho(\lambda)$ is the spectral reflectance of the object for which the tristimulus values are being calculated. The relative spectral irradiance distribution, $S(\lambda)$, represents the spectral characteristics of the illumination incident upon the object under the viewing conditions for which the tristimulus values are being calculated. The spectral tristimulus values (or color matching functions), $\bar{x}(\lambda)$, $\bar{y}(\lambda)$ and $\bar{z}(\lambda)$, show how much of each primary is required to match a monochromatic stimuli. The $\bar{x}(\lambda)$, $\bar{y}(\lambda)$ and $\bar{z}(\lambda)$ functions, based upon experimental data from many normal observers (Wright, 1928 and Guild, 1931), are printed in tabular form in many authoritative texts on colorimetry (e.g. Driscoll and Vaughan, 1978). The normalizing factor, k in Equations 1 - 3, is defined as

$$k = 100 / \left[\sum_{\lambda=380}^{780} S(\lambda) \bar{y}(\lambda) \Delta\lambda \right]. \quad (5)$$

Color Video Information

Color video signals are generally transmitted in one of two formats, composite video or separate RGB video signals. In keeping with Grassman's first law both of these formats are tridimensional in nature. The composite video format is the most commonly available format in video equipment since it is the format used by the television industry.

Encoding color information. The National Television Systems Committees' (NTSC) composite color video signal allows both color and monochrome monitors to receive the same signal. The Y signal, which contains the gray scale information, is combined with two amplitude modulated chrominance signals (I, in phase portion, and Q, quadrature, or shifted in phase by 90° , portion) to form the composite video signal. The three signal components (YIQ) are encoded in a band-sharing operation in which the chrominance signals are transmitted as a pair of sidebands having a common frequency of 3.58 MHz (Benson, 1986).

The intensity and chrominance information from a solid state color video camera outputting composite video is commonly derived from RGB information measured using appropriate optical interference filters and image sensors. Unfortunately, there is more than one "standard" definition of the RGB primaries. Research conducted in the U.S. using off-the-shelf color video equipment to obtain color images is based upon the NTSC standard for composite video and the RGB values used are those defined by the FCC (Federal

Communications Commission). The rotation matrix transforming the FCC RGB color space into the YIQ system is (Keil, 1983)

$$\begin{bmatrix} Y \\ I \\ Q \end{bmatrix} = \begin{bmatrix} 0.299 & 0.587 & 0.114 \\ 0.596 & -0.275 & -0.321 \\ 0.212 & -0.523 & 0.311 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \quad (6)$$

Decoding color information. When a camera, that produces color composite video, is used for color imaging the video signal must be decoded to access the color information. The most common technique of implementing color image processing is to digitize each of the RGB video signals separately, which gives three digital images for each primary color. Each pixel in the scene being analyzed is actually stored as a triplet of RGB values. The following rotation matrix can be used to transform the YIQ values from a composite color video signal into the FCC RGB color system (Benson, 1986)

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} 1.000 & 0.956 & 0.620 \\ 1.000 & -0.272 & -0.647 \\ 1.000 & -1.108 & 1.705 \end{bmatrix} \begin{bmatrix} Y \\ I \\ Q \end{bmatrix} \quad (7)$$

In addition to the RGB system the YIQ information can be transformed into the YØP color system. The hue and saturation values are simply the polar coordinate version of

the I and Q values and are determined by

$$\theta = \tan^{-1}(Q/I) \quad \text{and} \quad (8)$$

$$P = (I^2 + Q^2)^{1/2} \quad . \quad (9)$$

An overview of color imaging system that uses a solid state camera with color composite video output and an RGB video signal decoder is shown in Figure 3. In this example separate red, green, and blue interference filters are shown for simplicity but some video cameras derive the RGB information using overlapping patterns of different filters. The camera encodes the RGB information in composite video format using Equation 6. Analysis of the color information requires decoding into RGB color space using Equation 7. The color information, now in the computer system, can be transformed into YOP color space using Equations 6, 8 and 9.

Color Machine Vision

Several researchers have investigated the use of color, realizing that color information from natural scenes could greatly simplify the computer vision process. Konishi et al. (1984) used separate RGB optical filters and a black and white video camera to extract color information from a scene with color marked wires. An empirical relation using linear combinations of RGB intensity values was derived for each of

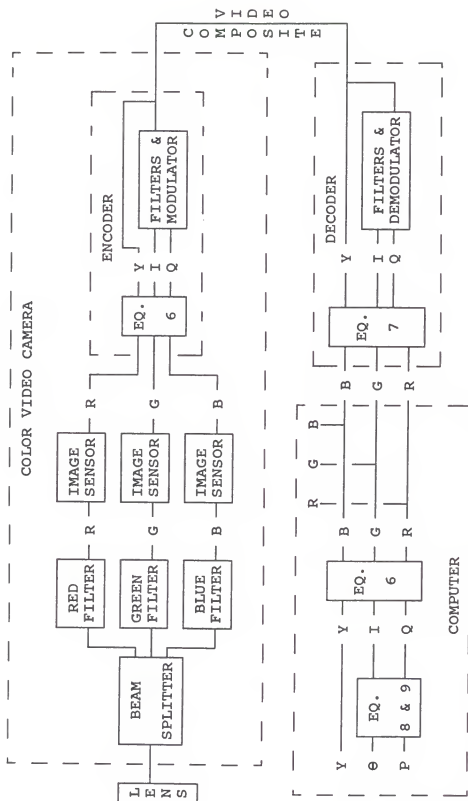


Figure 3. Color imaging system with a solid state color video camera, RGB decoder and computer.

four colors of wire. Threshold criteria for each of the empirical equations were used to distinguish the desired wire color. The system worked fairly well when a limited number of colors of wire were used. Solinsky (1985) used the chrominance information in a three dimensional scene for edge detection, reducing the computational complexity associated with edge detection using gray scale information. A black and white video camera and separate RGB optical filters were used to obtain RGB color images of the scene. Using Equations 6, 8 and 9 the color information was transformed in the YOP domain which was used instead of the RGB domain as the sensor space. Yoshimoto and Torige (1983) developed a high speed color information processing system for robot task control. With this system the color of an object was specified rather than the object's shape, greatly reducing the complexity of computations and making real-time control of a robotic manipulator feasible. A composite video camera with a RGB video encoder was used to record color information. Three simple comparators were used to classify the color of each element in the image into one of eight colors (black, blue, cyan, green, magenta, red, white and yellow). The system was capable of processing an image, to locate a colored object, in 50 ms (20 Hz) which was considered fast enough for real-time manipulator control. In addition to resolving the color information into only eight colors, the system had problems correctly classifying color correctly when the brightness of the image changed.

Keil (1983) developed a color vision system based upon a chromakeyer, a device that produces an analog TV signal which is brightest where the hue in the scene is in a selected range. The chromakeyer does not distinguish between colors which are equidistant from the selected hue and thus is "color blind" in these regions. For example, if orange was the selected hue, this system would not be able to distinguish between red and yellow. The recent development of low cost RGB cameras has made this technique less attractive especially if true color vision is desired.

Ohta (1985) developed a region analyzer for outdoor natural color scenes, in which the qualities of different methods for describing color (e.g. RGB, Y₀P etc.) were evaluated. The color feature set that performed the best seemed to depend upon the type of scene being segmented. Ohta found that in trying to completely segment a wide variety of outdoor natural scenes (e.g. from human portraits, to landscape scenes, to close-ups of cars) that intensity information, not color, was the most important feature, but that the quality of segmentation was often degraded by omitting color features. Because of the great diversity in the types of natural scenes studied, a rule based expert system was used to assist in the segmentation process, and although complex, the system often produced impressive results.

In three dimensional color space, Chrominance is defined (Jay, 1984) as a vector that lies in a plane of

constant luminance, and in that plane it may be resolved into components called chrominance components. In the YOP color system θ and P represent the chrominance components. Hue uniformities in colored scenes can be exploited for image segmentation allowing simpler and faster algorithms then might be possible if color were ignored. Hue can be used in the identification of objects under non-uniform illumination because hue is independent from the intensity of illumination reflecting from the scene. Kelley and Faedo (1985) used color vision for discrimination of color coded parts. They concluded that the phase-magnitude (i.e. hue-saturation) representation of the chrominance plane leads to computationally efficient scalar segmentation algorithms, and that saturated colors could be segmented using only hue and saturation information whereas nonsaturated colors (e.g. pastels and gray shades) require brightness information in addition to the hue and saturation information. Jarvis (1982) used color vision and a laser range finder to interpret three dimensional color scenes in an attempt to simplify the complex computations involved in three dimensional analysis. Hue could be used, because of its independence from intensity information in the scene, to identify those pixels belonging to the same connected component. Use of a laser range finder has the advantage, in addition to speed, of not being subject to the missing part problem encountered in binocular vision techniques. The missing part problem occurs when there is discrepancy

between the two images used to compute the range information due to the occlusion of one of the objects in one image and not in the other.

PROCEDURE

This chapter begins with a procedural outline of the research. Schematic diagrams of a robotic fruit harvester with a color vision system for real-time guidance and an aperture control system are presented. The equipment used to conduct this research is described and the chapter concludes with a description of the implementation of the color vision research.

Overview of Research

Quantifying Color Information in Natural Orange Grove Scenes

In order to quantify the color information present in natural orange grove scenes, the reflectance spectra of various objects in these scenes were measured. The perceived color of each object was quantified by calculating the tristimulus values (XYZ) from the spectrophotometric data using Equations 1 - 5. The tristimulus values were then transformed into the FCC RGB values using the following relation (Benson, 1986)

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = n \begin{bmatrix} 1.910 & -0.533 & -0.288 \\ -0.985 & 2.000 & -0.028 \\ 0.058 & -0.118 & 0.896 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \quad (10)$$

The color information was then transformed from the RGB color system to the Y Θ P color system using Equations 6, 8 and 9.

Once the colors of typical objects in a natural orange grove scene were quantified in each of the three color systems, the merits of each system were examined to determine which system had the most potential for using color vision information in robotic guidance. The XYZ and RGB color systems were considered to have the same potential for color imaging because both systems describe color as the addition of three primaries. The major difference between XYZ and RGB is that RGB primaries are real in a physical sense and the XYZ primaries are imaginary. The RGB system was preferred to the XYZ system for color imaging because color video output was easier to obtain in RGB format. The XYZ system was used primarily as the liaison between the spectrophotometric data and the RGB and Y Θ P systems.

The Y Θ P system was used to study the feasibility of using color information to detect and locate oranges in natural orange grove scenes. There are two major advantages of the Y Θ P system over the RGB system. First, the Y Θ P system is similar to the human vision system in the perception of color. A particular color, such as orange, is easily and uniquely described in the Y Θ P system by simply specifying a range of hue and saturation values. In the RGB color system the color orange cannot be described by simply specifying allowable ranges of red, green and, blue values

due to interactions between the three primaries. To classify a color in the RGB system not only must the RGB values be in the proper range but they must also be in proper proportion to one another. Second, the intensity, Y , is independent of θ and P . This means that, in theory, color or chrominance information (θ and P) can be used for robotic guidance even in scenes with non-uniform illumination and that colored objects should be identifiable using only two parameters (θ and P) instead of three (RGB).

Color Imaging

Hue and saturation thresholding. Typical orange grove scenes were recorded and stored on diskette in digitized format. From the color information extracted from the spectrophotometric data the desired hue and saturation threshold values for an orange were estimated. This information was used to show that natural orange grove scenes could be segmented into regions of fruit and background using only hue and saturation information from the scene. The images were segmented using the following rule

If $((\theta_{\min} < \theta < \theta_{\max}) \text{ AND } (P_{\min} < P < P_{\max}))$
 then classify as orange, else classify as background.

The process for determining threshold values ($\theta_{\min}, \theta_{\max}, P_{\min}$ and, P_{\max}) was inherently stochastic, due to variations in

illumination, camera adjustments, items in the image (e.g. clouds and dead leaves) for which spectrophotometric data was unavailable and natural variations in color from orange to orange. A trial and error process for determining the threshold values was used to obtain acceptable results. This technique provided an adequate method for color segmenting images when individual threshold values were selected for each image (Slaughter and Harrell, in press); unfortunately the threshold values varied from image to image. A systematic method for specifying the boundary of a selected color region needed was required.

Statistical pattern classification. A multivariate statistical pattern classification technique based upon probability theory was selected as a potential method for systematically classifying oranges and background in natural scenes. This method relies upon Bayes' rule for estimating the a posteriori probability that a particular RGB triplet belongs to one of two possible classes (oranges or background). The Bayesian technique selected is a form of discriminant analysis. Discriminant analysis attempts to assign, with a low error rate, an observation, x , of unknown classification, to one of two (or more) distinct groups (Lachenbruch, 1975).

The Bayesian approach assigns an observation to the class with the largest a posteriori probability. The Bayes' classifier was selected because it minimizes the total expected error in classifying objects, and from a

statistical point of view represents the optimum measure of performance (Tou and Gonzalez, 1974). When the data are multivariate normal and the covariance matrices are quite different (as is the case for oranges and background) the optimum classifier is a quadratic discriminant function (Duda and Hart, 1973). Little work in discriminant analysis for population densities other than the normal or the multinomial has been done (Lachenbruch, 1975). The assumption of a multivariate normal distribution is not completely accurate in this case because the data have a finite range of possible values rather than infinite range. However, Miller (1985) successfully used a Bayesian decision model to classify lemons into different grades based upon a finite range of visual blemish readings.

Because this technique incorporates the interactions between the color components, it was thought that color segmentation could be accomplished directly from RGB information as well as from θ and P information. Using RGB information is preferred from a computational standpoint over θ and P because RGB can be obtained directly from an RGB video camera whereas θ and P must be calculated from RGB values.

Quantifying color segmented image quality. Although the quality of a segmented image is a difficult concept to quantify precisely, the performance of the color segmentation of natural orange grove scenes was relatively simple to observe. A systematic method of measuring the

quality of a segmented image was needed to evaluate the performance of the Bayesian classifier. The technique used to quantify the quality of a color segmented image was based upon the ultimate goal--to harvest oranges. From the standpoint of manipulator control, the accuracy of the estimate of the centroid was the most important measure of image quality. Estimates of the horizontal and vertical diameters and the area of the fruit are also important in discriminating noise from oranges and in helping to determine whether the orange blob in the image is a single fruit or multiple fruits clustered together.

Image quality was quantified by five parameters: aq , xcq , ycq , xdq and ydq . The area quality parameter (aq) was defined as the estimate of the area of the object as a percent of the true area. The centroid quality parameters (xcq and ycq) were defined as the absolute error in the estimate of the centroid as a percent of the true diameter in the horizontal and vertical directions respectively. The diameter quality parameters (xdq and ydq) were defined as the estimate of the diameter as a percent of the true diameter in the horizontal and vertical directions respectively.

The object's true area and centroid were defined as the area and centroid calculated by chain code techniques. Chain code is an efficient method of representing the boundary of an irregularly shaped object using line segments and lends itself to calculation of such parameters as area

and centroid (Ballard and Brown, 1982). These parameters were used to evaluate the quality of color segmented images which were segmented using the Bayes classifier.

Real-Time Vision for Robotic Guidance

The main goal of this research was to provide a vision system capable of locating oranges in an orange tree at a sufficient rate to provide feedback control of a robotic manipulator. Observations of the natural oscillations of oranges in an orange tree indicate that when disturbed (as by the wind) they oscillate at frequency of about 0.5 Hz to 1 Hz. If the manipulator is going to accurately track the fruit as it moves in to pick, it needs to have a closed-loop bandwidth approximately ten times greater than the frequency of the fruit, or about 5 Hz to 10 Hz. The vision system must then supply new location information of the orange to the control algorithm at a rate about ten times the closed-loop bandwidth of the manipulator or about 50 Hz to 100 Hz.

The standard mode of operation of a video camera is to produce a new frame at a rate of about 30 Hz which is not in this range. Fortunately, standard video frame is comprised of two separate fields in an interlaced format. One field consists of the odd lines of a frame and the other field contains the even lines. In this format the even and odd fields are displayed alternately and together make up the entire image frame. If each field is used as a separate image the data rate doubles to 60 Hz. The compromise of

going to this faster data rate is a loss in spatial resolution in the vertical direction.

Once a digital color image has been acquired it can be segmented into regions using chrominance information. A color lookup table is an effective tool for use in segmenting color images because all computations can be done in advance. Each unique color in the color space can be assigned to a location in the lookup table and the classification information (i.e. fruit or background), corresponding to that color, stored there. The status of a color in a scene can be rapidly determined, for segmentation of an image, by checking that color's corresponding location in the lookup table. The size of the lookup table required is proportional to the total number of colors possible in a digitized image. If desired, multiple color lookup tables can be constructed in advance for different ranges of hues and saturations and then a particular table selected at the time of classification. For guidance of a robotic orange harvester only two regions are of interest, oranges and background. In this case the color lookup table is used to classify each color pixel in the image, and if desired, blackening any color pixel with a false status indicating that a particular color is outside the desired range of hue and saturation.

If a search algorithm is to be implemented at a real-time rate of 60 Hz it must examine no more pixels than necessary in detecting and locating a fruit. A search

starting at the center of the image and spiraling outward in a rectangular grid pattern will find any circular orange object in the image if the maximum distance (diagonally) between the elements of the grid is no larger than the diameter of the object. Fruit, occluded from view by leaves or branches, may not appear as large as a non-occluded fruit nor is it always circular in appearance so the size of the grid must be adjusted based upon the relative costs of speed and missed fruit.

Once a possible target is detected, the centroid and horizontal and vertical diameters of the object can be determined using the iterative technique developed by Harrell et al. (1985). This technique is based upon successively finding the chord through the object that is a perpendicular bisector to the previous chord beginning with an initial horizontal chord through the grid point that was initially detected as a possible target. Usually, three cycles of finding horizontal and vertical chords is sufficient to estimate the centroid. The lengths of the last horizontal and vertical chords found are used as estimates of the true horizontal and vertical diameters of the object. The main disadvantage with this technique is that it is inaccurate for objects with holes inside (like a cross-section of a toroid). The level of the error is dependent on the size and location of the hole, with the largest errors occurring when the hole is near the centroid of the fruit.

Aperture Control

Image quality and illumination. Successful color segmentation depends upon starting with a high quality image of the orange grove scene. Since the ultimate goal is to harvest oranges, the quality of the orange image is much more important than the quality of the background image. Without the proper amount of light the quality of the image is degraded.

A video camera is designed to operate within a specified range of illumination levels (e.g. the Javelin camera used in this work had a minimum illumination requirement of 20 lx and a maximum illumination level of 100,000 lx). A vision system that is designed to operate in the orange grove must have some means of adjusting for changes in illumination. Assume that the color (in the RGB system) of an object of interest is made up of nine parts red, five parts green, and one part blue (the approximate proportions of RGB for the color orange). When the image becomes slightly overexposed the red portion of the signal from the object will saturate before the green and blue portions. This means that as the illumination level increases and the image of the object becomes overexposed, the ratio of RGB values coming from the object will change. At first the object will appear more and more yellow in color as the mixture approaches equal parts red and green; then as the green signal saturates the object will appear more and more white as the mixture approaches equal parts

red, green, and blue. Underexposure causes similar problems without enough light to properly stimulate the color sensors.

When either of these illumination situations occur the color information becomes distorted and in some cases the image has less information than a black and white image. To overcome this problem cameras have two main methods to control the amount of light striking the image sensor: aperture control and artificial illumination (e.g. stroboscopic lamp).

A new measure of image intensity. Traditional autoiris lenses use the average intensity over the entire image as a feedback signal to control the aperture setting. Unfortunately, oranges often occupy 10% or less of the image area in a typical grove scene. If the other 90% of the area is made up of background material that has a different illumination level (e.g. as when the fruit is backlit by a bright blue sky) the oranges will not be adequately illuminated. To overcome this problem an aperture control scheme was proposed that would use only the average intensity of the orange pixels as a feedback signal to the autoiris lens.

Equipment Overview

Robotic Fruit Harvester with Color Vision System

A schematic diagram of the overall concept of a robotic fruit harvester using color vision for feedback control is shown in Figure 4. As described by Harrell et al. (1985)




Figure 4. A spherical coordinate robot using color vision for manipulator guidance.

the camera is mounted in the end-effector near the picking mechanism. In order to pick a fruit, its image must be approximately in the center of the field of view. Once the location of the centroid of the fruit is determined, its offset from the center of the image is calculated and used as an error signal for the algorithm controlling the joint actuators of the robot. Not shown in Figure 4 is a range sensor used to determine the distance the end-effector must

extend to pick the fruit and a proximity sensor used to determine when to activate the picking mechanism.

A schematic diagram showing the main components of a color vision system for robotic fruit harvest is shown in Figure 5. The color video information is decoded from NTSC composite video format to RGB video signals. The three separate RGB signals are digitized into three RGB images in real-time. The computer searches the color image (which now consists of three RGB images) using a color lookup table to detect and locate an orange. Once located, the centroid and diameter information can be used to control robot motion. A high-performance 8/16/32-bit bus (VME), designed for industrial applications, was used to allow the computer to communicate with the frame grabbers.

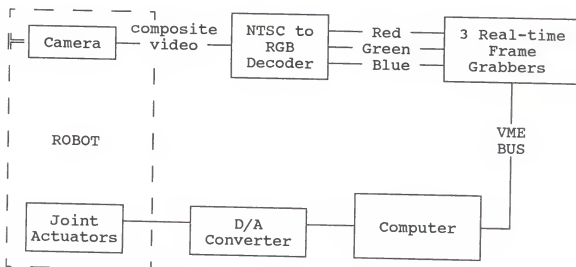


Figure 5. A color vision system for robot guidance.

Color Vision System with Object Oriented Aperture Control

A schematic diagram of a control system that will adjust the lens aperture of a camera using only the illumination information from objects of a specific color in the field of view is shown in Figure 6. Once an orange has been detected in a color image, the average intensity of the orange pixels are used as a feedback signal to the lens.




Figure 6. Schematic diagram of a color vision system with object oriented aperture control.

System Hardware

Spectrophotometric hardware. A computerized version of the spectrophotometer used by Norris and Butler (1961) with a detector head specially modified for the measurement of agricultural materials was used to make diffuse reflectance measurements. The spectrophotometer was located in the Instrumentation Research Laboratory, Sensors and Control Systems Institute, USDA-ARS, Beltsville, MD.

Color vision hardware. A Javelin (model 3012A) solid state color video camera and a Sony (model VO-5600) video tape recorder were used to tape orange grove scenes. The camera was equipped with a 50 mm focal length lens which had a aperture range of f1.4 to f22. A For-A (model DEC-100) NTSC color decoder was used to transform composite color video signals output from either the camera or the video tape player into RGB video signals.

The RGB video signals were input to three Datacube (model VVG-128) real-time video frame grabbers. Each frame grabber was setup to digitize and store in memory a 384 x 485 x 5 bit image frame. This spatial resolution (384 x 485) provides a total of 186,240 color pixels in each image. Every pixel in a color image consisted of a red, green, and blue triplet where each of the RGB values was digitized into one of 32 discrete intensity levels. Thirty-two discrete levels (five bits) of red, green and blue intensities were considered to be sufficient because experiments have shown that at a given light level, the human eye can only discern approximately 30 gray levels (Snyder, 1985).

A schematic diagram of a Datacube frame grabber is shown in Figure 7. Each of the three frame grabbers used in a true color imaging system occupy 256K bytes of memory which is byte addressable only. Each card has lookup tables on both input and output which allow simple preprocessing or pseudo color imaging, neither of which were implemented in this research.

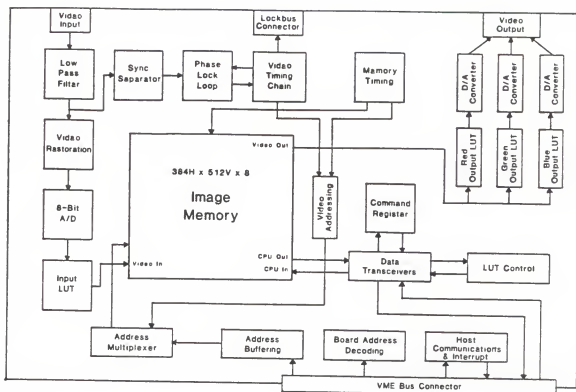


Figure 7. Datacube (model VVG-128) real-time video frame grabber (Datacube, 1985).

A Mizar computer system, with a Motorola 68020 microprocessor (running at 12.5 MHz), was used to process images. The 68020 microprocessor had 32 bit address and data registers which made 32 discrete levels of RGB values very convenient in the construction of color lookup tables used extensively in this work. The color space defined by 32 discrete levels of red, green, and blue intensities provided a total resolution of 32,768 colors. The computer had both hard and floppy disk drives on which digitized color images could be stored. A Sony (model PVM 1270Q) RGB

high resolution color monitor was used to display color images.

Aperture control hardware. A Sony CCD color video camera (model DXC-101) with a Cosmocar autoiris lens (8mm focal length) was used to implement the aperture control. The Sony camera was used in this portion of the research because its small size (154.5 mm x 67 mm x 63.5 mm) made it more compatible for location inside the robot arm, and the autoiris lens experiments required on-line computer control. The computer output the aperture control signal (in composite video format) via a Datacube video-frame grabber (model VVG-128) to the autoiris lens.

Implementation

Determining the Colors of a Typical Orange Grove Scene

Reflectance Spectra were collected for 'Valencia' orange peel, 'Valencia' tree leaves, 'Valencia' tree bark, orange grove soil, artificial plastic oranges and artificial plastic leaves. For all samples, except the soil, the spectrophotometer was setup to scan from 300 nm to 1,100 nm in 0.8 nm increments using a 3 mm slit width (10 nm bandwidth). The soil sample was scanned from 380 nm to 780 nm in 0.4 nm increments using a 2 mm slit width (7 nm bandwidth). The 300 nm to 1100 nm region of the spectrum is roughly equivalent to the region of sensitivity for silicon detectors used in solid state cameras. The 380 nm to 780 nm region of the spectrum is the visible region of the spectrum

and is used to calculate tristimulus values. Halon (polytetrafluoroethylene) powder was used as a reference material to calibrate the spectrophotometer because it is the reference material recommended by the U.S. National Bureau of Standards for spectral reflectance measurements (Weidner and Hsia, 1981).

The CIE tristimulus values were calculated from the spectral reflectance data, assuming the specimen was illuminated with standard illuminant C (daylight with a correlated color temperature of 6774K), with Equations 1 - 5. The tristimulus values were transformed into the FCC RGB values using the relation shown in Equation 10. The tristimulus values for an albedo sample (the whitish inner portion of citrus rind) were used as a reference for the intensity level of the sample when the XYZ values were converted to RGB values. The RGB values were normalized to simplify direct comparison between RGB values determined from spectrophotometric data and the RGB values observed in color images. The normalizing constant, n , was set equal to 0.444 so that the red intensity level (which was the largest of the RGB values) would be normalized to the maximum value (31) possible for the digital color imaging system. This normalizing constant was then used for all of the other samples so that the RGB and P values would lie in the same range of values as the color information from the digital color images. The XYZ, RGB and θP values were also calculated for the CIE standard illuminant C for comparison

purposes (using the relative spectral irradiance distribution data given by Driscoll and Vaughan, 1978). The RGB values for illuminant C were normalized independently of the selected grove items ($n = 0.310$) in order to maintain an adequate illumination level on the grove samples.

Determining the brightness, hue, and saturation (YOP), corresponding to a particular triplet of RGB values requires a coordinate transformation. The rotation matrix transforming the FCC RGB color space into the YIQ system was shown in Equation 6. Using Equations 8 and 9, the levels of hue and saturation were obtained for each sample.

Processing Color Images

Taping natural orange grove scenes. Images of natural scenes were collected from 'Valencia' orange groves near Lake Alfred, FL. Because transportation of the computer vision system was impractical, video tape recordings were made of the orange grove and then images were digitized from the video tape at a later time.

There were three major parameters that directly influenced the quality of the video image: focus, aperture, and white balance. These three parameters were set by eye in the grove using a color video monitor for visual feedback. White balance was used to adjust the camera for different spectral irradiances in the lighting (e.g. tungsten vs. daylight). The procedure used was to place a white object in front of the camera so that it filled the

entire field of view and then by adjusting the white balance the relative gains of each of the RGB signals were adjusted so that the intensity of the red signal is equal to the intensity of green and blue signals. The white balance was adjusted by eye using visual feedback from a RGB video monitor. Once the white balance was adjusted the setting was not changed for the video taping session so that in the event of the white balance not being perfectly adjusted any bias would be present at the same level in all images. Newer cameras are available with an auto white balance circuit that greatly simplifies adjustment of the white balance.

Developing color lookup tables. Once a digital color image has been acquired it can be segmented into regions using hue and saturation thresholds or a Bayesian classifier. Because of the time required to process each of the 186,240 RGB triplets in a color image, a lookup table was developed to determine which particular RGB triplets (out of the 32,768 possible) were within a specified color region. For this purpose a 32 kilobit (i.e. four kilobytes) color lookup table was constructed, where each of the 32,768 possible colors was represented by an individual bit in the lookup table. The lookup table was laid out into 1024 long words. Each long word was 32 bits in length and each bit represented the 32 levels of blue ($0 \leq B \leq 31$) possible. The red and green intensity levels were used to indicate the

address of the desired long word in the table. The address was calculated by

$$\text{Address} = 32 * R + G + A_s, \quad (11)$$

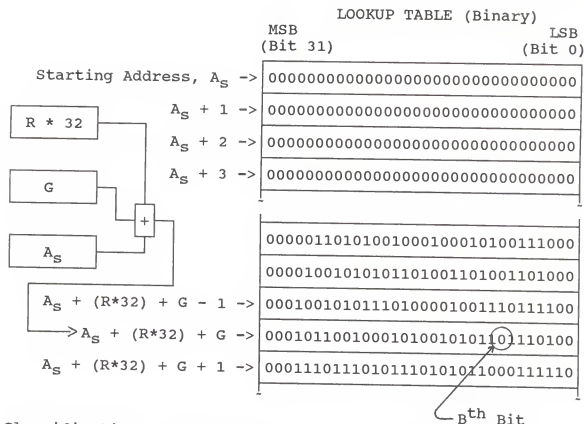
where $0 \leq R \leq 31$,

$0 \leq G \leq 31$ and

A_s is the starting address of the lookup table.

Once an address was calculated from the red and green values, the blue value was used to determine which bit in the desired long word needed to be examined. By its true (1) or false (0) status, the bit specified if the corresponding RGB triplet was within the desired color range. The process of classifying an RGB triplet using a color lookup table is illustrated in Figure 8.

In binary arithmetic multiplication by a power of 2 can be achieved by shifting the binary point in the same fashion as shifting the decimal point can represent multiplication by a power of 10 in base 10. Thirty-two is two raised to the fifth power; thus the multiplication in Equation 11 can be accomplished by shifting the binary point of the red value five times to the right. A five bit shift operation is approximately one order of magnitude faster in MC68020 assembly code than multiplying by 32. This feature represents a large reduction in the time required to lookup the status of an RGB triplet in the table.



Classification = State (True or False) of Bit, Selected using the R and G Values to Determine the Table Address and the B Value to Determine which Bit to Test.

As shown, the value of B would be 8 and the classification would be false (0), indicating background.

Figure 8. Illustration using R and G values to locate the desired lookup table address, then using the B value to determine which bit to use in classifying an RGB triplet.

This method of implementing the color segmentation algorithm required four memory accesses, a five bit shift operation, an addition, and a comparison to classify each color pixel triplet in the image. A color segmentation algorithm, using the color lookup table and written in assembly code, could classify approximately 75,000 color

pixels per second and an entire color image could be segmented in less than 2.5 seconds.

Statistical pattern classification. Implementation of a Bayes classifier requires knowledge of the probability density function characterizing each of the classes, in this case oranges and background.

The classifier assigns a RGB triplet to the class of oranges if

$$D_o(x) > D_b(x) \quad , \quad (12)$$

where D_o is the discriminant function for oranges and D_b is the discriminant function for the background. The random variable x is a vector containing θ and P values or RGB values. The discriminant function, D_i , derived for minimum error rate classification of class w_i is

$$D_i(x) = P(w_i|x) \quad . \quad (13)$$

The conditional probability, $P(w_i|x)$, is the probability that w_i is the true classification, given that the "color" of the pixel was x . Using Bayes' rule, $P(w_i|x)$ can be written

$$P(w_i|x) = \frac{p(x|w_i) P(w_i)}{\sum_j p(x|w_j) P(w_j)} \quad , \quad (14)$$

where

$$\sum_j p(x|w_j) P(w_j) = p(x) \quad . \quad (15)$$

The conditional density function of the measurement x is $p(x|w_i)$ and $P(w_i)$ is the a priori probability of class w_i . The density function, $p(x)$, is not a function of the class w_i and is usually dropped from the discriminant analysis in Equation 14 because it does not help in distinguishing between classes. For the multivariate normal case we assume the conditional densities $p(x|w_i)$ are multivariate normal (i.e. $p(x|w_i) \approx N(\mu_i, C_i)$, where μ_i is the true population mean vector and C_i the covariance matrix for the class w_i). Assuming that $C_o \neq C_b$ (the o subscript indicates oranges and the b subscript indicates the background) the simplified discriminant functions are

$$D_i(x) = -1/2[(x-\mu_i)^T C_i^{-1} (x-\mu_i) + \log|C_i|] + \log P(w_i). \quad (16)$$

In order to implement this technique a training data set must be used to provide estimates for μ_i and C_i . Once the means and covariances have been estimated a color lookup table can be built by calculating D_o and D_b for each of the 32,768 possible colors. If $D_o > D_b$ for a specific color, then the corresponding bit in the table is set otherwise it is cleared.

Two data sets, one for training and one for evaluation, were collected from an image of a natural grove scene to

compare the classification of pixels based on θ and P with a classification based on RGB information using Proc Discrim (SAS, 1985). The a priori probabilities for oranges and background were chosen to be proportional to the number of pixels from each class in the training data set for Proc Discrim. In addition this statistical procedure tests the covariance matrices to see if they are statistically different as assumed. The results (summarized in Table 1) show that a discriminant function based on RGB information performed as well as one based on θ and P . The analysis also showed that the homogeneity of the covariance matrices was rejected at the $\alpha = 0.10$ level of significance which was expected because the background was a much more diverse group of colors than the pixels from oranges.

Measuring color segmented image quality. The quality of color segmented images was quantified by five parameters, aq , xcq , ycq , xdq , and ydq . These five parameters were defined as

$$aq = 100 \times a/a' , \quad (17)$$

$$xcq = 100 \times |xc - xc'|/xd' , \quad (18)$$

$$ycq = 100 \times |yc - yc'|/yd' , \quad (19)$$

$$xdq = 100 \times xd/d' \text{ and,} \quad (20)$$

$$ydq = 100 \times yd/d' , \quad (21)$$

where

a = area of object in segmented image,

a' = area of object in original image,

xd = horizontal diameter of object in segmented image,
 xd' = horizontal diameter of object in original image,
 yd = vertical diameter of object in segmented image,
 yd' = vertical diameter of object in original image,
 xc = horizontal component of centroid of object in
segmented image,
 xc' = horizontal component of centroid of object in
original image,
 yc = vertical component of centroid of object in
segmented image and
 yc' = vertical component of centroid of object in
original image.

The area and centroid were calculated using chain code techniques. The chain code for determining the "true" parameter values (a' , xc' , yc' , xd' and yd') from the unsegmented image was constructed by manually tracing the boundary of the object. The chain code for objects in a segmented image were constructed automatically assuming the object was four-connected (Ballard and Brown, 1982). Any time clustered fruit appeared in an image such that the fruit were touching or one overlapped the other, they were treated as one object.

Table 1. Evaluation of Statistical Classification Method

Classification using Hue and Saturation

Training Data Set

<u>Type of Pixels</u>	<u>Total</u>	<u>Classified As</u>	
		<u>Orange</u>	<u>Background</u>
Orange	220	203	17
Background	813	80	733

Evaluation Data Set

<u>Type of Pixels</u>	<u>Total</u>	<u>Classified As</u>	
		<u>Orange</u>	<u>Background</u>
Orange	252	201	51
Background	1147	36	1111

Classification using Red, Green, and BlueTraining Data Set

<u>Type of Pixels</u>	<u>Total</u>	<u>Classified As</u>	
		<u>Orange</u>	<u>Background</u>
Orange	220	201	19
Background	813	0	813

Evaluation Data Set

<u>Type of Pixels</u>	<u>Total</u>	<u>Classified As</u>	
		<u>Orange</u>	<u>Background</u>
Orange	252	206	46
Background	1147	18	1129

=====

Evaluating color segmentation. The quality of color segmentation using the Bayesian statistical classification technique was evaluated using images of natural orange grove scenes. Fourteen images were collected for evaluation of the color segmentation technique. The 14 images were selected to represent the variety of natural scenes observed in a typical 'Valencia' orange grove. To insure consistent quality of the images used for evaluation, the intensity of each image was adjusted prior to digitization so that the average intensity, \bar{V}_O , of the oranges in the image was within the range

$$13 < \bar{V}_O < 18 \quad . \quad (22)$$

This range of intensity was selected because it is in the middle of the dynamic range of possible intensity values for this system. Manual aperture control was used in the grove to approximately set the intensity within the desired range and then the intensity level was adjusted (if necessary) in the laboratory using the gain adjustment on the For-A decoder.

One of the 14 images was used as a training image to estimate the means and covariance matrices. The criterion for selecting a training image was based on the requirement that the background objects in the image typify the backgrounds expected to be encountered in actual harvesting conditions. The image used to estimate the object and

background population means and covariances is shown in Figure 9.

The number of pixels used from the training image for estimating the statistical parameters was based upon consideration for the time required for the operator to classify pixels as to orange or background and upon the quality of the segmentation of the training image. The data for each of the four data sets were collected by recording the RGB values at each of the grid points of four different grid sizes and querying the operator if each grid point was part of an orange or part of the background. A separate color lookup table was created from the mean and covariance matrices determined from each of the four data sets using Equation 16. The a priori probabilities required by the Bayesian classifier were assumed to be 0.10 for oranges and 0.90 for the background. In cases where the a posteriori probabilities were nearly equal the choice of a priori probabilities could have an effect on the final classification. In general, the a priori probability of finding an orange ranged from 5% to 25% depending upon the distance from the fruit to the camera and upon the number of fruit in the field of view. A conservative estimate for the a priori probability of finding an orange was used because it was considered better to fail to harvest a fruit rather than risk having the robot try and pick a tree limb possibly damaging both robot and tree.



(a)



(b)

Figure 9. Color image of orange grove scene used to train the Bayesian Classifier. (a) Digitized orange grove scene; (b) Color segmented image, using 20x25 data.

The image quality, using the measures of image quality previously defined, for four different training data set sizes is shown in Table 2. The grid size of 20 x 25 was selected as the best compromise between quality of segmentation and time required for the operator to classify the pixels in the training set. The color lookup table created from the training image was used to color segment the remaining 13 images of grove scenes. The quality of the segmentation process was measured in the same manner as previously defined.

Table 2. Image quality vs. size of training data set

Grid Size	Orange Pixels	Background Pixels	XCQ %	YCQ %	XDQ %	YDQ %	AQ %
55x70	6	43	3.09	9.06	79.63	84.67	67.59
35x45	21	97	2.47	7.32	88.89	90.59	75.39
20x25	71	261	2.47	7.32	88.89	90.94	77.45
10x15	222	842	1.23	6.62	88.89	90.59	78.85

Evaluating Real-Time Orange Location

As stated previously the color segmentation algorithm developed can segment pixels in an image at a rate of about 75,000 color pixels per second. If a search algorithm is to be implemented at a real-time rate of 60Hz, it can only search through 1250 color pixels per cycle. A grid spacing of 20 pixels in horizontal direction by 25 pixels in the

vertical direction was used in this research to locate oranges in an image. For a grid of this size there are approximately 400 grid points in the image to check, accounting for one third of the time allowed per cycle in the worst case. Once an orange object is detected the centroid and diameters are measured using the iterative technique developed by Harrell et al. (1985). When determining the centroid and diameters of a possible target, the spatial resolution was increased by decreasing the grid size of pixels examined from 20 x 25 to 4 x 4 resulting in an image 121 pixels in the horizontal direction by 96 pixels in the vertical direction.

The real-time search algorithm was evaluated using the same lookup table created with the Bayesian classifier where the image, shown in Figure 9, was used for training. The time required to locate the orange closest to the center of the image and determine its centroid and diameters was measured. The 13 images used to evaluate the quality of the color segmentation process were used to benchmark the real-time algorithm and the parameters xcq , ycq , xdq , and ydq were also determined for comparison to those found with the chain code technique.

Aperture Control Experiments

Lens aperture--image intensity relationships. A schematic diagram of a control system, shown in Figure 6, adjusts the lens aperture of a camera using only the

illumination information from objects of a specific color in the field of view. Once an orange has been detected in a color image, the average intensity of the orange pixels is used as a feedback signal to the lens.

To evaluate the potential of such a scheme a series of images was collected under varying aperture settings and two different illumination conditions. Artificial plastic oranges were placed in the foliage of shrubbery on the campus of the University of Florida, Gainesville, to simulate a natural orange tree scene. Images were collected under two different illumination conditions, in direct sunshine on a sunny day and, on an overcast day. Two color lookup tables were created, one for each illumination condition, using θ and P thresholding for simplicity as described previously. Each image was segmented using the appropriate lookup table, the quality of the image segmentation was measured as described previously, and the average intensity of the pixels classified as oranges in the image were recorded.

Characterizing autoiris capabilities. A commercially available solid state camera with an autoiris lens was evaluated to determine its suitability for implementing the proposed object oriented aperture control concept. The autoiris lens used was typical of commercially available autoiris lenses used in surveillance applications. This type of lens uses the standard composite video signal coming from the camera for feedback. The real-time orange location

algorithm was enhanced to implement the proposed aperture control technique based upon the intensity of the orange pixels in the image. Once an orange was found in the image, the computer calculated the average intensity of the fruit as the centroid and diameters were determined. The computer output the average intensity of the fruit to the lens using a Datacube video-frame grabber. This method of controlling the lens aperture was used because the lens required a composite video signal as input and the frame grabber could be easily setup to output a composite video signal in real-time.

An open-loop step test was used to try and mathematically model the lens and camera. The open loop response to a step change in the input video signal, shown in Figure 10, suggests that the system might be approximated by a first order system with an integrator. The system was found to be non-linear by observing the open-loop step test of the lens-camera system under varying magnitudes of step size. Non-linearity due to both saturation and a delay were observed. Because of the non-linearity, the system parameters were determined on-line using closed loop proportional control as shown in Figure 11. Proportional control was selected because it is simple to implement and is fairly robust. The control system was inherently digital with an analog plant; the zero order hold (ZOH), shown in Figure 11, represents the Datacube card. The lens-camera system was modeled using an open-loop transfer function for

a first order system with an integrator (shown in Laplace transform style). The two constants, K and T_i , are the open-loop steady state gain and the first-order time constant respectively. The sampling rate of the system is represented by T . The intensity setpoint is Y_{set} and the average intensity of orange pixels is \bar{Y}_O . A series of closed-loop step tests was performed by stepping Y_{set} from 3 to 26, for varying proportional gains, K_p .




Figure 10. Open-loop step test results of autoiris lens.




Figure 11. Closed-loop Proportional Control System for Aperture Control.

RESULTS

This chapter begins with a summary of the color characteristics of typical objects in a natural orange grove scene. Thirteen color images of orange grove scenes and the corresponding color segmented (using the Bayesian classification technique) images are presented. The quality of the segmentation process is summarized for both the chain code technique and the real-time algorithm for fruit location. The chapter concludes with the results from the aperture control experiments.

Color Characteristics of Typical Orange Grove Objects

The diffuse spectral reflectances of typical objects from an orange grove environment are shown in Figures 12 through 24. Figure 25 shows the spectral irradiance relative to equal energy white of the CIE standard illuminant C (the assumed illuminant for determining the tristimulus values). Figures 26 and 27 show the diffuse spectral reflectances of an artificial plastic orange and leaf respectively, which were used in the aperture control experiments. The reflectance curves are plotted as the logarithm (base 10) of the reflectance versus wavelength.




Figure 12. Diffuse spectral reflectance from the albedo of an orange.




Figure 13. Diffuse spectral reflectance from the first sample of orange peel.




Figure 14. Diffuse spectral reflectance from the second sample of orange peel.




Figure 15. Diffuse spectral reflectance from the third sample of orange peel.




Figure 16. Diffuse spectral reflectance from the fourth sample of orange peel.




Figure 17. Diffuse spectral reflectance from an orange peel with slight regreening.




Figure 18. Diffuse spectral reflectance from a regreened orange peel.




Figure 19. Diffuse spectral reflectance from orange tree bark.




Figure 20. Diffuse spectral reflectance from a medium green orange tree leaf, top.




Figure 21. Diffuse spectral reflectance from a medium green orange tree leaf, bottom.




Figure 22. Diffuse spectral reflectance from a dark green orange tree leaf, top.




Figure 23. Diffuse spectral reflectance from a dark green orange tree leaf, bottom.




Figure 24. Diffuse spectral reflectance from a sample of orange grove soil.




Figure 25. Spectral irradiance for CIE standard illuminant C (Driscoll and Vaughan, 1978).




Figure 26. Diffuse spectral reflectance from an artificial plastic orange.




Figure 27. Diffuse spectral reflectance from an artificial plastic leaf.

The color characteristics of typical objects from natural orange grove scenes are summarized in Table 3. While by no means an exhaustive sampling, the data presented in Table 3 reinforces the hypothesis that color information in orange grove scenes can be used to differentiate oranges from other objects. Examination of the θ chrominance parameter provides the clearest description of the difference in color between objects.

Color Segmentation of Natural Orange Grove Scenes

The color segmentation process, using a color lookup table created with a statistical pattern classification technique that attempts to minimize the error of misclassification, was applied to thirteen digital color images from natural orange grove scenes. The sample means and covariance matrices used to build the color lookup table are shown in Table 4 (data from the image shown in Figure 9).

Table 4. Estimates of orange and background color parameters

	Orange				Background			
	<u>Mean</u>	<u>Covariance</u>			<u>Mean</u>	<u>Covariance</u>		
Red	21.2	19.9	11.0	4.7	10.7	90.0	78.7	64.4
Green	13.8	11.0	11.9	10.7	10.8	78.7	70.5	58.4
Blue	6.0	4.7	10.7	14.7	8.7	64.4	58.4	51.6

Table 3. Typical color parameters of oranges and background objects


SAMPLE ID.	TRISTIMULUS VALUES			RGB VALUES			CHROMINANCE
	X	Y	Z	R	G	B	
ALBEDO (INSIDE ORANGE PEEL)	63.58	66.37	56.21	31.00	30.41	20.51	-39.95 4.60
ORANGE PEEL (SAMPLE 1)	31.41	27.31	6.79	19.29	10.43	2.08	-5.15 7.99
ORANGE PEEL (SAMPLE 2)	30.77	27.93	6.56	18.63	11.26	1.94	-10.24 7.51
ORANGE PEEL (SAMPLE 3)	23.29	21.02	6.33	13.95	8.40	2.01	-8.56 5.42
ORANGE PEEL (SAMPLE 4)	37.25	33.21	6.68	22.86	13.11	1.88	-8.61 9.52
ORANGE SLIGHT REGREENING	30.10	30.22	7.59	17.39	13.57	2.21	-24.71 6.52
REGREENED ORANGE PEEL	20.16	22.58	7.84	10.75	11.12	2.45	-47.33 3.78
CITRUS TREE BARK	9.52	9.81	6.46	4.93	4.46	2.30	-30.54 1.13
MEDIUM GREEN CITRUS LEAF, TOP	7.94	10.15	5.28	3.65	5.47	1.77	-86.09 1.54
MEDIUM GREEN CITRUS LEAF, BOTTOM	21.26	25.14	14.63	10.20	12.83	5.05	-72.65 3.12
DARK GREEN CITRUS LEAF, TOP	6.19	7.31	5.80	2.78	3.70	2.08	-92.70 0.70
DARK GREEN CITRUS LEAF, BOTTOM	14.86	17.57	10.79	7.06	8.96	3.75	-74.95 2.09
ORANGE GROVE SOIL (SANDY)	21.30	21.47	21.80	10.19	9.47	8.09	-17.65 0.91
CIE STANDARD ILLUMINANT C	98.07	100.00	118.25	30.96	31.00	30.92	-86.61 0.05
ARTIFICIAL ORANGE	30.94	21.93	7.28	20.11	5.84	2.54	11.79 9.77
ARTIFICIAL LEAF	8.06	10.30	6.59	3.55	5.54	2.29	-95.57 1.44

The thirteen original images and the results of color segmentation are shown in Figures 28 through 40. The quality of the color segmentation process was quantified by the five parameters a_q , xc_q , yc_q , xd_q , yd_q and is documented in Table 5.

In some of the images (e.g. Figures 30, 31, 32, 37 and, 40) more than one closed orange region was present in the segmented image, for these images the quality of the segmentation of each separate orange region was quantified separately. Each digital color image consisted of 384×485 pixels, the upper left corner of the image was defined as the origin, $(0,0)$, and the lower right corner of the image was defined to have the coordinates $(384, 485)$. Using this coordinate system, the location of the centroid (xc, yc) of each orange region evaluated was specified in Table 5 for future reference. The orange shown in the upper right corner of Figure 31 was partially occluded by tree branches and thus divided into many smaller regions upon segmentation; only the largest portion located at $(37, 51)$ was evaluated for segmentation quality.




(a)




(b)

Figure 28. Color image of orange grove scene with oranges and background leaves. (a) Digitized orange grove scene; (b) Color segmented image.



(a)




(b)

Figure 29. Color image of orange grove scene with oranges and background leaves. (a) Digitized orange grove scene; (b) Color segmented image.




(a)




(b)

Figure 30. Color image of orange grove scene with oranges, background leaves and branches. (a) Digitized orange grove scene; (b) Color segmented image.




(a)



(b)

Figure 31. Color image of orange grove scene with oranges, background leaves and branches. (a) Digitized orange grove scene; (b) Color segmented image.



(a)



(b)

Figure 32. Color image of orange grove scene with oranges, background leaves and sky. (a) Digitized orange grove scene; (b) Color segmented image.



(a)



(b)

Figure 33. Color image of orange grove scene with oranges, background leaves and sky. (a) Digitized orange grove scene; (b) Color segmented image.



(a)



(b)

Figure 34. Color image of orange grove scene with oranges, background leaves, sand and sky. (a) Digitized orange grove scene; (b) Color segmented image.




(a)




(b)

Figure 35. Color image of grove scene with an orange, background leaves and sand. (a) Digitized orange grove scene; (b) Color segmented image.



(a)



(b)

Figure 36. Color image of grove scene with an orange, background leaves and sky. (a) Digitized orange grove scene; (b) Color segmented image.




(a)



(b)

Figure 37. Color image of orange grove scene with oranges, background leaves and sand. (a) Digitized orange grove scene; (b) Color segmented image.



(a)



(b)

Figure 38. Color image of orange grove scene with oranges, background leaves and sky. (a) Digitized orange grove scene; (b) Color segmented image.



(a)




(b)

Figure 39. Color image of grove scene with an orange, background leaves and sand. (a) Digitized orange grove scene; (b) Color segmented image.



(a)



(b)

Figure 40. Color image of orange grove scene with oranges and background leaves. (a) Digitized orange grove scene; (b) Color segmented image.

The results in Table 5 show that on the average the estimate of horizontal centroid is within 5% of the horizontal diameter to the "true" centroid and similarly the estimate of the vertical centroid is within 6% of the vertical diameter to the "true" centroid. For an orange 8 cm in diameter this means that the estimated centroids would be within approximately ± 5 mm of the "true" centroid of the fruit under conditions similar to those used in this research.

Table 5. Quality of Color Segmentation

Figure	Location		\bar{Y}_O	Quality				
	XC	YC		AQ %	XCQ %	YCQ %	XDQ %	YDQ %
28	259	266	16	87.1	3.13	7.65	105.4	98.2
29	187	219	16	86.5	4.97	0.48	97.2	104.8
30	313	185	15	88.5	1.33	7.69	105.3	84.1
30	190	137	15	89.5	0.94	8.10	75.5	93.2
31	37	51	15	104.7	2.56	2.08	105.1	106.2
31	243	299	15	89.9	2.73	1.10	96.4	87.3
32	246	227	16	59.5	4.76	10.20	76.2	75.5
32	223	130	16	66.2	9.86	4.95	88.7	89.3
32	167	222	16	63.9	8.82	4.91	82.4	82.8
33	198	243	17	62.0	9.09	1.62	77.8	58.9
34	245	256	16	91.9	0.56	3.20	101.1	95.5
35	220	212	16	81.9	2.46	8.69	100.0	92.3
36	150	229	16	81.7	6.06	6.21	73.5	95.9
37	121	240	15	67.1	7.08	3.64	69.0	109.4
37	319	245	15	37.0	5.88	12.00	54.9	50.0
38	169	179	15	65.5	0.96	2.59	69.2	84.5
39	224	229	15	76.9	3.31	10.10	102.4	79.4
40	104	278	17	73.1	1.10	6.94	80.2	91.0
40	260	188	17	72.9	4.30	11.80	108.6	79.9
40	194	335	17	94.7	5.13	3.06	92.3	107.1
Ave.			16	77.0	4.25	5.85	88.1	88.3

Note: \bar{Y}_O is the average intensity of the orange pixels in the image.

required to determine the centroid and diameters was 10.8 ms which is less than the maximum time allowed per cycle of 16.7 ms.

Aperture Control

The resultant variation in image quality with image intensity is shown in Figures 41 through 45. These five figures show the five quality parameters plotted against the average intensity (\bar{V}_O) of the orange pixels from each of the color segmented images at each combination of aperture setting and illumination condition. The results show that the quality of the color segmented image is best when the average intensity of the orange pixels is near the middle of the dynamic range of possible intensity levels. Thus, to maximize image quality, the setpoint, V_{set} , used in the control system shown in Figure 11 should be approximately 16.

Figures 46 to 50 show closed-loop step responses of the system with different levels of gain, K_p . The closed-loop system shows a transient response similar to a standard second-order system. Assuming the sampling rate is fast enough that the system can be approximated by an analog control system, the closed-loop transfer function for a second-order system can be written as

$$G(s) = \frac{\omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \quad (23)$$




Figure 41. Plot of area quality versus average intensity of orange pixels.




Figure 42. Plot of centroid quality (horizontal component) versus average intensity of orange pixels.




Figure 43. Plot of centroid quality (vertical component) versus average intensity of orange pixels.




Figure 44. Plot of diameter quality (horizontal) versus average intensity of orange pixels.




Figure 45. Plot of diameter quality (vertical) versus average intensity of orange pixels.

The damping ratio, ζ , is the ratio of the actual damping in the system to the level of damping when the system is critically damped (Ogata, 1970). The undamped natural frequency of the system is defined as ω_n . The performance characteristics of a control system can often be determined from the transient response to a step input (Ogata, 1970). The performance characteristics can be determined by measuring any two of several transient response characteristics. For underdamped systems, two of the easiest characteristics to measure are the peak time, t_p , and the maximum percent overshoot, M_p . The peak time is the time required to reach the first peak of the overshoot. The maximum percent overshoot is the amount of maximum overshoot output by the system expressed as a percent of the steady-state output of the system.

A second-order system should have a damping ratio in the range

$$0.4 < \zeta < 0.8,$$

for a desirable transient response (Ogata, 1970). If the damping ratio is less than 0.4, the amount of overshoot may be excessive and if the damping ratio is greater than 0.8, the system may respond sluggishly. Too much overshoot or a sluggish response can cause the vision system to saturate or blacken the entire image destroying vital data for the robot




Figure 46. Closed-loop step test with a proportional gain, $K_p = 1$.




Figure 47. Closed-loop step test with a proportional gain, $K_p = 2$.




Figure 48. Closed-loop step test with a proportional gain, $K_p = 3$.




Figure 49. Closed-loop step test with a proportional gain, $K_p = 4$.




Figure 50. Closed-loop step test with a proportional gain, $K_p = 5$.

guidance system. It is also desirable to have as high a closed-loop bandwidth as possible so that the system can respond quickly to disturbances in the illumination level.

Examination of the closed-loop response curves shown in Figures 46 to 50 indicated that the transient response of Figure 47 ($K_p = 2$) most closely represented the desired transient response. However, further examination revealed that although the maximum overshoot and peak time were as expected the rise times of the curves in Figures 48 to 50 were nearly identical to the rise time of Figure 47. (For underdamped second-order systems the rise time is often defined as the time required for the response to rise from

0% to 100% of its final value.) This indicated that the system was velocity saturated and that implementing a more sophisticated controller would not improve the speed of response.

The closed-loop bandwidth, ω_b , of the autoiris lens was estimated from the maximum percent overshoot and peak time in Figure 47. The maximum percent overshoot was estimated as

$$M_p = \frac{(27 - 3) - (26 - 3)}{(26 - 3)} = 0.043 \quad .$$

The maximum percent overshoot can be written (Ogata, 1970),

$$M_p = \frac{-\pi(\zeta/(1 - \zeta^2)^{1/2})}{e} \quad . \quad (24)$$

The damping ratio can be calculated using Equation 24 which gives

$$\zeta = 0.71 \quad .$$

The peak time was estimated as $t_p = 0.45$ s. The peak time can be written (Ogata, 1970)

$$t_p = \frac{\pi}{\omega_n(1 - \zeta^2)^{1/2}} \quad . \quad (25)$$

The natural frequency can be determined using Equation 25 and the damping ratio which gives

$$\omega_n = 9.9 \text{ r/s} = 1.6 \text{ Hz} \quad .$$

The closed-loop bandwidth is commonly defined to be the frequency at which the magnitude of the closed-loop frequency is 3 dB below its zero frequency value (Ogata, 1970). The closed-loop bandwidth can be determined by

$$|G(j\omega_b)|_{\text{dB}} = -3 \text{ dB},$$

where $G(s)$ was defined in Equation 23 and $s = j\omega_b$. When $\zeta = 0.71$, as in this example, $\omega_b = \omega_n$ which gives

$$\omega_b = 1.6 \text{ Hz}.$$

This estimate of the closed-loop bandwidth must be used with caution because it was derived using a linear model for a system that displayed non-linear characteristics. This analysis illustrated the constraints of using a traditional autoiris lens with a real-time vision system for robot guidance.

CONCLUSIONS

This research demonstrated that color in natural orange grove scenes provides sufficient information about the scene to allow a vision system based solely on color to detect and locate oranges. The color of oranges differs from orange grove background items (e.g. leaves, branches, soil and sky) sufficiently to allow color images to be segmented into regions of fruit and background using color information.

The time constraints on a vision system, designed for guiding a robotic manipulator in the harvest of fruit, were determined from estimates of the natural frequency of oscillation of fruit on an orange tree. These parameters indicate that the vision system would need to detect and locate oranges in the field of view at a rate of 50 Hz to 100 Hz. Standard color video signals could supply information at this rate (60 Hz) if each field of the interlaced video format were used as a separate video image. The vision system would then have 16.7 ms in which to search each image for oranges to provide real-time vision feedback for the robot arm.

Three color systems (XYZ, RGB and Y Θ P) were studied as to their suitability for use in segmenting color images. The Y Θ P system was best suited to color segmentation from a

user standpoint because Y θ P is the most similar to the way humans perceive color. In addition, only two parameters (θ and P) were required to adequately define the desired color region in the Y θ P system whereas three parameters were required in the other systems. From an implementation standpoint the RGB color system was preferable because it is commonly used in color video cameras for sensing color information. Also, if a RGB camera was used, less information degradation would occur, less hardware would be required and less time would be required because no coordinate system translation would be done.

Although θ and P thresholding proved feasible, a multivariate statistical classification technique based upon Bayesian probability theory was investigated. This technique was advantageous because it provided a systematic method for describing the desired region in color space. Due to the incorporation of interactions between parameters, the Bayesian classifier performed equally well using either θ and P or RGB information. Because of advantages in implementation, a Bayesian classifier based upon RGB information was used to evaluate the capabilities of the color segmentation process.

Due to the vagueness in the concept of image quality, quantifiable quality parameters were developed that described the quality of a segmented image in terms pertinent to robotic fruit harvesting. These quality parameters were based upon the ability to accurately

estimate the centroid and size of the fruit in the image. The results show that, on the average, the quality of the color segmented image was such that the estimated centroid of the oranges in the image differed, on the average, from the true centroid by $\pm 6\%$ of the diameter of the fruit. The area of the fruit in the color segmented image was, on the average, 77% of the area of the fruit in the unprocessed image.

The real-time aspect of fruit detection and location was evaluated using a real-time search algorithm developed by Harrell et al. (1985). Using this technique the system required 4.07 ms on the average to detect and locate an orange in the image. In the worst case 10.8 ms was required using this technique which was within the allotted 16.7 ms. Due to the iterative nature of this technique some degradation in the estimate of the fruit location was experienced. The estimated centroid differed, on the average, from the true centroid by $\pm 10\%$ of the diameter of the fruit.

Scene illumination and lens aperture setting significantly affected image quality. Results showed that the quality of color segmented images was best when the average intensity of the orange pixels was in the middle of the dynamic range of possible intensity values. These results are consistent with typical aperture control settings for autoiris lenses, the main difference being in the source of the intensity information. Typically autoiris

lenses use the entire field of view (or some subset) for feedback, in the case of robotic fruit harvest only the intensity of the fruit is of interest. By maintaining an object oriented aperture control (in this case fruit oriented) maximum image quality should be maintained.

The dynamic response of a typical autoiris lens was evaluated to determine its suitability for a real-time vision guidance system. Closed-loop analysis showed that the lens' response time was less than desired for a real-time vision system. To maximize system performance, the closed-loop bandwidth of the autoiris lens should be at least as high as that of the robotic manipulator.

Highly non-uniform illumination across a single fruit presents a problem which cannot be solved by aperture control. Figures 36, 37 and 38 show examples of non-uniform illumination caused when both diffuse and direct sunlight illuminate different parts of the same fruit. In these three examples aperture control adjusted the intensity so that the majority of the fruit was correctly illuminated but not the whole fruit. One possible solution to this problem would be the use of supplemental illumination (e.g. stroboscopic lamp). Strobe light is commonly used by photographers to illuminate shadows of non-uniformly illuminated subjects and might be used to make the surface of the fruit more uniformly illuminated. Another possible solution would be to use some type of translucent material, as an artificial canopy, to diffuse the sunlight.

To achieve optimal performance in segmenting color images the Bayesian classifier should be retrained whenever the appearance of objects in the scene change. Although data from a single image were successfully used to train the classifier in this research, data from more than one image may be required in an actual harvesting situation. In an actual orange harvesting operation, varietal and seasonal changes in fruit appearance as well as changes in background can be expected. Ideally, some type of adaptive or self-tuning technique is desired to continuously update the training data set and retrain the Bayesian classifier when the performance falls below an acceptable level.

Two other areas for future work deal with the problems of clustered fruit and with fruit partially occluded from view. Both of these problems were beyond the scope of this research but need to be addressed if successful fruit harvest is to be accomplished. In the case of occlusion, the problem is to adjust the estimate of the centroid location to compensate for an unknown portion of the fruit which is not visible. Shape is the most logical parameter to use in this case due to the spherical nature of the fruit. Clustered fruit must also be distinguished but shape alone is insufficient due to the occlusion problem. In this case, real-world size is the most logical parameter to use. Real-world size could be estimated from vision information if the distance from the camera to the fruit was also known.

REFERENCES

- Ballard, D.H., and C.M. Brown. 1982. Computer vision. Prentice-Hall, Inc. Englewood Cliffs, NJ.
- Benson, K.B. 1986. Television engineering handbook. McGraw-Hill. New York.
- Brown, G.K., C.E. Schertz, and C.K. Huszar. 1971. Fruit-bearing characteristics of orange and grapefruit trees in California. ARS #42-181 U.S. Dept. of Ag. Washington, D.C.
- Committee on colorimetry, Optical Society of America. 1944. The psychophysics of color. J. Opt. Soc. Am. 34(5):245-266.
- Coppock, G.E. 1977. Citrus harvesting in Florida. Proc. Int. Soc. Citriculture 2:393-397.
- Coppock, G.E. 1983. Robotic principles in the selective harvest of Valencia oranges. Proc. of the First International Conf. on Robotics and Intelligent Machines in Agriculture. ASAE. St. Joseph, MI. pp. 138-145.
- Coppock, G.E., and P.J. Jutras. 1960. An investigation of the mobile picker's platform approach to partial mechanization of citrus picking. Florida State Hort. Soc. 73:258-263.
- Datacube, Inc. 1985. VVG-182 video acquisition and display board user's manual. Datacube Inc. Peabody, MA.
- Driscoll, W.G., and W. Vaughan. 1978. Handbook of optics. McGraw-Hill. New York.
- Duda, R.O., and P.E. Hart. 1973. Pattern classification and scene analysis. John Wiley and Sons Inc. New York.
- Gaffney, J.J. 1969. Reflectance properties of citrus fruit. Trans. of the ASAE 16(2):310-314.

- Grand d'Esnon, A. 1984. Robotic harvesting of apples. Proc. of the First International Conference on Robotics and Intelligent Machines in Agriculture. Oct. 2-4, 1983. ASAE. St. Joseph, MI. pp. 112-113.
- Grand d'Esnon, A. 1985. Robotic harvesting of apples. Proc. of the Agri-Mation 1 Conf. and Expo. Feb. 25-28, 1985, Chicago, IL. pp. 210-214.
- Grassman, H. 1853. On the theory of compound colors. Phil. Mag. Ser. 4 7:254-264 and plate III.
- Guild, J. 1931. The colorimetric properties of the spectrum. Phil. Trans. Roy. Soc. London. [A], 230:149.
- Guralnik, D.B. 1982. Webster's new world dictionary. Simon and Schuster. New York.
- Hackwood, S., and G. Beni. 1984. Sensor and high-precision research. Robotics research. MIT Press. Cambridge, MA. pp. 529-545.
- Harmon, L.D. 1982. Automated tactile sensing. International Journal of Robotics Research. 1(2):3-32.
- Harrell, R.C. In press. Economic analysis of robotic citrus harvesting. Trans. of the ASAE.
- Harrell, R.C., P.D. Adsit, and D.C. Slaughter. 1985. Real-time vision-servoing of a robotic tree fruit harvester. ASAE Paper #85-3550, St. Joseph, MI.
- Hollaender, A. 1956. Radiation biology. Vol. III. McGraw-Hill. New York.
- Jarvis, R.A. 1982. Expedient 3-D robot colour vision. Proc. of the 2nd International Conference on Robot Vision and Sensory Controls. November 2-4, 1982. Stuttgart, Germany. pp. 327-338.
- Jay, F. 1984. IEEE standard dictionary of electrical and electronics terms. IEEE Inc. New York.
- Judd, D.B. 1933. The 1931 CIE standard observer and coordinate system for colorimetry. J. Opt. Soc. Am. 23:359-374.
- Judd, D.B. 1950. Colorimetry. National bureau of standards circular 478. Washington, D.C.
- Katsushi, I., B.K.P. Horn, S. Nagata, T. Callahan, and O. Feingold. 1984. Picking up an object from a pile of objects. Robotics research. MIT Press. Cambridge, MA. pp. 139-162.

- Keil, R.E. 1983. Machine vision with color detection. Proc. of SPIE. Intelligent Robots: Third International Conference on Robot Vision and Sensory Controls RoViSec3. Part II. Nov. 7-10, 1983. pp. 503-512.
- Kelley, R.B. and W. Faedo. 1985. A first look into color vision. Proc. of SPIE. Intelligent Robots and Computer Vision. 579:96-103.
- Konishi, T., M. Takagi, and J. Kitsuki. 1984. Shape reconstruction of wires using color images for automatic soldering. Robotics research. MIT Press. Cambridge, MA. pp. 389-399.
- Lachenbruch, P.A. 1975. Discriminant analysis. Hafner Press. New York.
- Martin, P. L. 1983. Labor-intensive agriculture. Scientific American 249(4):54-60.
- Miller, W.M. 1985. Decision model for computer-based grade separation of fresh produce. Trans. of the ASAE 28(4):1341-1345.
- Moravec, H.P. 1984. Locomotion, vision, and intelligence. Robotics research. MIT Press. Cambridge, MA. pp. 215-224.
- Nakagawa, Y. and T. Ninomiya. 1984. The structured light method for inspection of solder joints and assembly robot vision systems. Robotics research. MIT Press. Cambridge, MA. pp. 355-369.
- Nassau, K. 1980. The causes of color. Scientific American 243(4):124-153.
- Newton, I. 1730. Opticks. G. Bell and Sons. Reprinted by Dover Publications. New York. 1952.
- Norris, K.H., and W.L. Butler. 1961. Techniques for obtaining absorption spectra on intact biological samples. IRE transactions on Bio-Medical Electronics BME-8(3):153-157.
- Ogata, K. 1970. Modern Control Engineering. Prentice-Hall, Inc. Englewood Cliffs, NJ.
- Ohta, Y. 1985. Knowledge-based interpretation of outdoor natural color scenes. Pitman Publishing Inc. Marshfield, MA.
- Overheim, R.D., and D.L. Wagner. 1982. Light and color. John Wiley & Sons. New York.

- Parrish, E.A. and A.K. Goksel. 1977. Pictorial pattern recognition applied to fruit harvesting. Trans. of the ASAE 20(5):822-827.
- Pejsa, J.H., and J.E. Orrock. 1983. Intelligent robot systems: potential agricultural applications. Proc. of the First International Conf. on Robotics and Intelligent Machines in Agriculture. ASAE. St. Joseph, MI. pp. 104-111.
- Pinson, L.J. 1983. Robot vision: an evaluation of imaging sensors. Proc. of the SPIE - Robotics and Robot Sensing Systems. 442:15-26.
- Rosenfeld, A., and A.C. Kak. 1982. Digital picture processing. Academic Press. Orlando, FL.
- SAS, Institute Inc. 1985. SAS user's guide: statistics. Version 5 edition. SAS Institute Inc. Cary, NC.
- Schertz, C.E., and G.K. Brown. 1968. Basic considerations in mechanizing citrus harvest. Trans. of the ASAE 11(2):343-346.
- Slaughter, D.C., R.C. Harrell, P.D. Adsit, and T.A. Pool. 1986. Image enhancement in robotic fruit harvesting. ASAE Paper #86-3041. St. Joseph, MI.
- Slaughter, D.C., and R.C. Harrell. In press. Color vision in robotic fruit harvesting. Trans. of the ASAE.
- Snyder, W.E. 1985. Industrial robots: computer interfacing and control. Prentice-Hall Inc. Englewood Cliffs, NJ
- Solinsky, J.C. 1985. The use of color in machine edge detection. Proc. Vision'85 Conference, March 25-28, 1985. Detroit, MI. Machine Vision Association of SME. pp. (4-34)-(4-52).
- Tou, J.T., and R.C. Gonzalez. 1974. Pattern recognition principles. Addison-Wesley Pub. Co. Reading, MA.
- Tuttle, E.G. 1983. Image controlled robotics in agricultural environments. Proc. of the First International Conference on Robotics and Intelligent Machines in Agriculture. Oct. 2-4, 1983. ASAE. St. Joseph, MI. pp. 84-95.
- Weidner, V.R., and J.J. Hsia. 1981. Reflection properties of pressed polytetrafluoroethylene powder. J. Opt. Soc. Am. 71(7):856-861.
- Whitney, J.D. 1977. Mechanical removal of fruit from citrus trees. Proc. Int. Soc. Citriculture 2:407-412.

- Whittaker, A.D., G.E. Miles, O.R. Mitchell, and L.D. Gaultney. 1984. Fruit location in a partially occluded image. ASAE Paper #84-5511. St. Joseph, MI.
- Wolfe, R.R., and M. Swaminathan. 1986. Determining orientation and shape of bell peppers by machine vision. ASAE Paper #86-3045. St. Joseph, MI.
- Wright, W.D. 1928. A redetermination of the trichromatic coefficients of the spectral colors. Trans. Opt. Soc. 30:141.
- Yoshimoto, K. and A. Torige. 1983. Development of a colour information processing system for robot vision. Developments in robotics. Edited by B. Rooks. pp. 161-166. Bedford, IFS Publications LTD. North Holland Publ. Co. Amsterdam.

BIOGRAPHICAL SKETCH

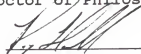
David Charles Slaughter was born January 18, 1960, in Tucson, AZ. He received his elementary education from the Flowing Wells Unified School District in Tucson and graduated from the Special Projects High School for advanced study in Tucson in June of 1978.

The author graduated, with high honors, in June of 1982 from the University of California, Davis, with a Bachelor of Science degree (agricultural engineering). Upon graduation he was awarded the departmental citation for outstanding undergraduate achievement in agricultural engineering.

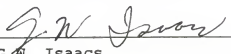
In August of 1984 the author completed the requirements for the degree of Master of Science (biological and agricultural engineering with a minor in mathematics) at North Carolina State University, Raleigh.

In August of 1984 the author accepted a position as research engineer with the United States Department of Agriculture, Agricultural Research Service (ARS). At that time he became a participant in the ARS Research Engineer Recruitment and Development Program and enrolled in the Graduate School of the University of Florida, Gainesville, for the degree of Doctor of Philosophy.

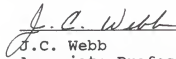
I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.


R.C. Harrell, Chairman
Assistant Professor of
Agricultural Engineering

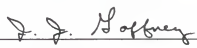
I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.


G.W. Isaacs
Professor of
Agricultural Engineering


I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.


J.C. Webb
Associate Professor of
Agricultural Engineering

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.


J.J. Gaffney
Associate Professor of
Agricultural Engineering

I certify that I have read this study and that in my opinion it conforms to acceptable standards of scholarly presentation and is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.


T.E. Bullock
Professor of
Electrical Engineering

This dissertation was submitted to the Graduate Faculty of the College of Engineering and to the Graduate School and was accepted as partial fulfillment of the requirements for the degree of Doctor of Philosophy.

August 1987

Hubert A. Bevis

Dean, College of Engineering

Dean, Graduate School